Guest Editorial Special Section on Real-Time and (Networked) Embedded Systems—II

THE design of embedded systems requires the development of efficient platforms supported by innovative technologies, methodologies, and algorithms. The interaction with the environment, the timing of the operations, the communication networks, and last but not least the market requirements, however, pose significant challenges to the designers. This is especially true for safety-critical applications, which require guarantees in the face of complex distributed implementations and of the intricacies of the interaction between the components of the system. New methods and techniques are, therefore, being developed by researchers in academia and industry alike, to more easily and efficiently take advantage of the opportunities afforded by ubiquitous computing and connectivity to offer new and original functions and solutions, reaching far into unconventional areas of application, promising new solutions to problems such as energy conservation, transportation, environment preservation, and health care.

This Special Section presents some of the most significant recent research work representing the state-of-the-art in the area of the embedded and/or networked real-time systems for automation systems. The section builds on top of and extends contributions to the track on "Real-Time and (Networked) Embedded Systems," which was part of the IEEE International Conference on Emerging Technology and Factory Automation (ETFA), held in Palma de Mallorca, Spain, in September 2009. The section is divided into three groups of papers that relate to each other topic-wise. The first group contains papers dealing with the fusion of control theory and scheduling of real-time systems, and the second group deals with timing analysis of real-time networks. Finally, the third group of papers investigates design, modeling, and timing analysis of embedded real-time systems.

A. Real-Time and Control

Real-time systems theory often builds upon a set of assumptions of the system and its corresponding model. However, inherent in the complex nature of interactions between the real-time system components and its environment, it is hard to predict and model exact characteristics of these using classical techniques. Instead, combining control theory with real-time scheduling mechanisms introduces flexibility in both the modeled system and its implementation, allowing the system to adapt to changing behaviors and scenarios, and thereby opening up for adoption of real-time research in a whole new domain of applications.

Digital Object Identifier 10.1109/TII.2010.2078931

The papers presented in this first group deal with several applications of control theory in the context of real-time systems. The first paper presents a framework constructed with two nested controllers establishing Quality-of-Service (QoS) on resource and application level, respectively, for real-time applications. The second paper introduces the concept of anytime computing in the context of control of embedded real-time systems. The third paper investigates how spare bandwidth can be used to increase control performance of Controller Area Network (CAN)-based networked embedded control systems. In the following, the three papers in this group are summarized with more details.

The paper "On the Integration of Application Level and Resource Level QoS Control for Real-Time Applications" presents a resource reservation framework for dynamic QoS in the context of soft real-time systems. There are many examples of applications with soft real-time requirements, e.g., multimedia, video and image processing, etc., where "soft" simply means that it is not catastrophic as some timing requirements occasionally would be violated. In order to make good usage of the resources (e.g., CPU, memory, power) in the context where these applications are executed, it is often not a good decision to make resource reservations based on applications' worst-case resource consumptions. Such a design would require too many resources, hence making the application or product too expensive. Instead, allowing for a more flexible resource reservation behavior, introducing control mechanisms where resource allocation parameters can be dynamically changed during runtime, an efficient usage of resources is made while at the same time fulfilling a sufficient level of QoS.

The paper introduces a resource reservation framework, in which the resource allocation parameters are dynamically changed online using two nested control feedback loops. The inner loop controls resource-level QoS, whereas the outer loop controls application-level QoS. The presented framework has been thoroughly evaluated in a Linux implementation.

The paper "Design of Embedded Controllers Based on Anytime Computing" presents the application of anytime control concepts in the context of embedded real-time systems. Inherent in the ever increasing complexity of embedded systems, where a number of applications often must share a common hardware with limited resources, there is a need to efficiently utilize these resources. During runtime, an application may experience changing availability of resources over time, typically due to interference from other applications. It is not a feasible solution to design the embedded system according to worst-case application resource requirements, but instead applications should adjust their resource parameters according to the current resource availability, and the current application needs. It is here that control theory comes into play. However, classical control design methodologies suffer from important limitations, e.g., conservative parameter settings. It is the thesis of this paper that these limitations may be resolved or relaxed by using unconventional (given the context of embedded real-time systems) control theory.

This paper introduces a design methodology for how the anytime control paradigm can be used in real-time systems. Central to the approach presented in the paper is that a control law is split into a set of subroutine calls, for which a stochastic modeling framework has been developed and evaluated. The framework has been evaluated in the context of real-time task systems, and experimental results suggest that the presented approach performs well in resource constrained embedded controllers.

The paper "Run-Time Allocation of Optional Control Jobs to a Set of CAN-Based Networked Control Systems" presents an approach to increase control performance of networked control systems built with the Controller Area Network (CAN). CAN is a commonly used network technology for networked control systems, and it has a particular feature of predictable message transmissions inherent in a priority-based message arbitration mechanism. Therefore, CAN is a popular choice for many networked embedded applications requiring real-time guarantees and/or good control performance.

The paper introduces a new way to utilize the unused bandwidth of a CAN-based networked control system, to increase control performance. Such systems are traditionally constructed through periodic transmission of control messages. By transmitting additional aperiodic messages within the unused bandwidth, i.e., the bandwidth not used by the periodic control messages, control performance can be significantly improved. Hence, the presented approach uses the properties of CAN to transmit as many control actions as possible to achieve the best control performance. This technique has been evaluated and validated through a proof-of-concept implementation.

B. *Real-Time Networks*

Central to any networked embedded system with real-time requirements is that the underlying network technology can provide real-time guarantees for message transmissions. In this Special Section, we have two papers on real-time networks. Both of these papers present research results on switched Ethernet and the application of network calculus for timing analysis in such systems. The first paper presents a new more accurate approach to analyze such systems, and the second paper presents a particular case of application and analysis of switched Ethernet technology in avionics. In the following, these two papers are presented with more details:

The paper "Improving the Worst-Case Delay Analysis of an AFDX Network Using an Optimized Trajectory Approach" presents how network calculus can be improved in the analysis of switched Ethernet for avionics applications. The ability to perform accurate end-to-end timing predictions is instrumental in the certification process of avionics systems, and in analyzing such end-to-end temporal behavior, accurate analysis techniques are imperative. Today, network calculus is the dominating approach for such analysis of AFDX networks, however, network calculus often results in pessimistic timing predictions.

The contribution of this paper is the optimization of a trajectory approach in the context of network calculus for worst-case end-to-end temporal delay computation of AFDX networks. The resulting new analysis suggests a significant improvement (10%) of analytical upper bounds of packet delays. This improvement has the potential to simplify certification of AFDX networks, and it allows for better utilization of the network resource in avionic applications.

The paper "Performance Analysis of a Master/Slave Switched Ethernet for Military Embedded Applications" presents an application of Flexible Time-Triggered Switched Ethernet (FTT-SE) in the context of avionics networks. Older generation military avionic communication networks are not sufficient in meeting the requirements of the next-generation military embedded applications. Therefore, a new networking technology must be introduced in this context, and looking at avionic networks in general, Switched Ethernet is an obvious candidate. However, in supporting an easy migration of existing military subsystems to a new switched Ethernet architecture, it is beneficial to keep their current centralized communications scheme by using a master/slave transmission control also in the new architecture, such as the one provided by FTT-SE.

The paper presents the application and usage of FTT-SE in the context of military avionics applications, intended at allowing for the migration of existing applications to the new network technology. Moreover, appropriate network calculus-based analysis techniques have been developed for determining the temporal characteristics of such applications. The results suggest that the proposed approach (technology + analysis) is a viable solution for the next generation of military avionics applications.

C. *Embedded Real-Time Systems*

The practical implementation of real-time embedded systems must deal with a number of issues related to the interaction with the environment, the use of resources, and the precise scheduling of the operations. The three papers in this group approach these problems from different angles. The first paper proposes a common but customizable and extensible hardware platform that includes services to manage robustness, efficiency, and complexity. The second paper presents an extension to a well-known modeling formalism to account for timing and resources. Finally, the third paper approaches the scheduling problem to determine an optimal pattern of preemptions to minimize the scheduling overhead. In the following, these three papers are summarized in more detail.

The paper "A Cross-Domain Multiprocessor System-on-a-Chip for Embedded Real-Time Systems" addresses the problem of designing architectures for embedded systems that can be applied to various domains of application. Establishing common platforms in this area is important to avoid costly redesign cycles and to take advantage of multidomain reuse. In their contribution, the authors identify several challenges for the design of embedded architectures, including complexity, robustness, and energy efficiency. They present GENESYS, a multiprocessor system on a chip which can be customized at different levels to individual applications. The architecture is composed of a domain-independent hardware core that provides the basic architectural services. These can be extended through system components and through specialized services for complexity management and robustness. In turn, different implementation options for these services allow the designer to customize the architecture to optimally trade off the parameters of interest. The component-based architecture is real-time aware and communicates with the environment through the exchange of messages. The architecture design is demonstrated on a case study taken from the automotive industry, which shows that heterogeneous services, such as control and multimedia, can be integrated with different criticality levels.

The paper "Timed and Resource-Oriented Statecharts for Embedded Software," emphasizes that handling resources and time in the design of embedded systems is an important factor that must be accounted for at the early stages of the modeling level. In many systems, the interaction with the environment must be based on precise timing relationships, which are significantly affected by the amount of resources available in the implementation. Traditional functional models, however, are often limited to expressing data and control relationships only, while they offer well established verification methods.

The paper takes a model-based approach and describes an extension to the well-known Statecharts formalism to model and analyze timed and resource-limited embedded systems. The resulting model is called Timed and Resource-Oriented Statecharts (TRoS). TRoS is based on Algebra for Communicating Shared Resources (ACSR) and extends Statecharts by adding the notion of "timed action node," modeling the scheduling and the use of resources of the different components of the system. In this formalism, I/O communication is instantaneous, while the use of resources takes one or more time units, and it is subject to resource availability. The paper presents the syntax and the semantics of TRoS, and shows its applicability to a case study coming from railway control systems. Various timing properties are also verified using the VERSA tool.

The paper "Limited Preemption EDF Scheduling of Sporadic Task Systems" presents timing analysis of real-time task systems following the limited preemption task model. It is well known that fully preemptive scheduling of independent tasks scheduled according to the Earliest Deadline First (EDF) scheduling policy is optimal with respect to schedulable resource utilization, given task scheduling in a single processor real-time system. However, when it comes to practically realize such scheduling of tasks, the actual scheduling mechanism and the runtime preemption of tasks will consume CPU resources manifested as runtime overhead. In fact, lowering the number of preemptions is desirable in order to reduce scheduling overhead, as each time a task is preempted the scheduler is involved and will consume CPU resources. On the other hand, having no preemptions at all introduces overheads in terms of tasks blocking each other, e.g., a low-priority task may block a high-priority task due to the fact that preemption is not allowed. Hence, there is a tradeoff centering around the number of preemptions that should be allowed given a specific system configuration, and when these preemptions should be made in order to minimize overall runtime overheads.

The paper presents results given the limited preemption task model. These results are of particular interest when it comes to implementing an EDF scheduler on a real platform.

A special section like this one relies on the active support of several people, and we would like to take the opportunity to thank them all: the authors for their contributions and their cooperation in promptly replying to the reviewers' comments; the reviewers, for their thorough reviews and comprehensive comments, which contributed significantly to the quality level of the papers that were accepted for publication; and finally, Prof. Richard Zurawski, Editor-in-Chief, for his guidance in preparing and finalizing this Special Section, from the very first steps until its publication.

THOMAS NOLTE, *Guest Editor*
Mälardalen University
MRTC
Västerås 721 23, Sweden
thomas.nolte@mdh.se

ROBERTO PASSERONE, *Guest Editor*
University of Trento
Department of Information Engineering
and Computer Science
Trento I-38122, Italy
roberto.passerone@unitn.it

Thomas Nolte (S'01–M'09) received the B.Eng., M.Sc., Licentiate, and Ph.D. degrees in computer engineering from Mälardalen University (MDH), Västerås, Sweden, in 2001, 2002, 2003, and 2006, respectively.

He was a Visiting Researcher at the University of California, Irvine (UCI), Los Angeles, in 2002, and a Visiting Researcher at the University of Catania, Catania, Italy, in 2005. He was a Postdoctoral Researcher at the University of Catania in 2006, and at MDH in 2006–2007. He became an Assistant Professor at MDH in 2008. He is now an Associate Professor of Computer Science at MDH since 2009. His research interests include predictable execution of embedded systems, design, modeling and analysis of real-time systems, multicore systems, distributed embedded real-time systems.

Dr. Nolte is Vice Chair of the IEEE Industrial Electronics Society (IES) Technical Committee on Factory Automation (TCFA), since 2009, President of The Swedish National Real-Time Association (SNART), since 2008, and Program Leader of the PROGRESS National Strategic Research Centre at Mälardalen Real-Time Research Centre (MRTC), since 2007. He is a Member of the Editorial Board of Elsevier's Journal of Systems Architecture: Embedded Software Design, since 2009, Guest Editor of several sections in the IEEE TRANSACTIONS ON INDUSTRIAL INFORMATICS, General Co-Chair of the IEEE International Symposium on Industrial Embedded Systems (SIES), 2011, Co-Organizer of international events including the International Workshop on Compositional Theory and Technology for Real-Time Embedded Systems (CRTS), 2008, 2009 and 2010 [co-located with IEEE Real-Time Systems Symposium (RTSS)], and Program Co-Chair of the IEEE International Conference on Emerging Technologies and Factory Automation (ETFA), the Real-Time and (Networked) Embedded Systems track, 2008, 2009 and 2010, and the IEEE International Workshop on Factory Communication Systems (WFCS), 2008 (Work-in-Progress Track) and 2012.

Roberto Passerone (S'96–M'05) received the Laurea degree (*summa cum laude*) in electrical engineering from the Politecnico di Torino, Torino, Italy, in 1994, and the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, in 1997 and 2004, respectively.

From 1998 to 2005, he was with Cadence Design Systems, Berkeley, CA, where he held various positions from Senior Member of Technical Staff in the System Level Design Product Group, to Research Scientist in the Cadence Berkeley Laboratories. Since 2006, he has been an Assistant Professor with the Department of Information Engineering and Computer Science at the University of Trento, Trento, Italy. At the University of Trento, he is the Leader of the COMBEST European Project, and of the Artist Design Network of Excellence. His research interests include the design and implementation of high-performance microprocessor, system level design, communication design, and formal methods.

Dr. Passerone has been Co-Chair of the Semantics Working Group and the Rosetta Working Group in the Accellera Standardization Committee and was Secretary of the IEEE P1699 Rosetta Standard. He is also Co-Chair of the Embedded Systems Subcommittee of the IEEE IES TCFA. He has been Guest Editor for the IEEE TRANSACTIONS ON INDUSTRIAL INFORMATICS in 2008 and 2009 and co-organizer of several events, including the Real-Time and (Networked) Embedded Systems Track at the IEEE International Conference on Emerging Technologies and Factory Automation in 2008, 2009 and 2010, and the IEEE Symposium on Industrial Embedded Systems in 2009 and 2010.